

Conversations among Competitors

By JEREMY C. STEIN*

Conversation is a central part of economic life. A wide range of information gets passed from one person to another via word-of-mouth communication, and a number of authors argue that the nature of connections between people—i.e., the structure of social networks—can have a profound influence on how far and how efficiently information spreads across the economy.¹ For example, in the context of financial markets, Robert J. Shiller (2000, 155) writes: “Word-of-mouth transmission of ideas appears to be an important contributor to day-to-day or hour-to-hour stock market fluctuations”

In many sorts of conversations, it can be taken for granted that the participants will communicate honestly with one another. If one friend asks another for her opinion of, say, a restaurant or movie, it is hard to see why the response would be anything other than completely truthful. There are also a number of important cases, however, where it is less obvious that an honest exchange of information can be expected. Often the participants in a conversation are competitors with one another, at least along some dimension, which means that when one player’s information set is improved, the other player may be made worse off. Consider the following illustrations.

Professional Money Managers.—Harrison Hong, Jeffrey D. Kubik, and Stein (2005) document that mutual-fund managers in a given city tend to have trading behavior that covaries more strongly with other managers in the same city, as opposed to with managers in different cities. Lauren H. Cohen, Andrea Frazzini, and Christopher J. Malloy (2007) uncover a similar correlation structure in the trades of mutual-fund managers who went to college together. These findings can be interpreted as evidence of word-of-mouth communication among those fund managers who are socially connected to each other.² But given that professional money managers have strong incentives to care about relative performance, it is not clear why one would tell another honestly about an attractive trading opportunity that he has discovered.³

Knowledge Spillovers in Silicon Valley.—Many observers argue that the free flow of ideas across firms has been a key factor in fostering the high rate of technological progress in the Silicon Valley. For example, AnnaLee Saxenian (1994, 2–3) writes:

Silicon Valley has a regional network-based industrial system that promotes collective learning and flexible adjustment among specialist producers of a complex of related technologies. . . . *Companies compete intensely while at the same time learning from one another about changing markets and technologies through informal communications*

(Italics added)

* Department of Economics, Harvard University, Cambridge, MA 02138 and National Bureau of Economic Research (e-mail: jeremy_stein@harvard.edu). I am grateful to the National Science Foundation for financial support, and to Sam Hanson for outstanding research assistance. Thanks also to Oliver Hart, seminar participants at Dartmouth University, and two anonymous referees for helpful comments.

¹ See, e.g., Glenn D. Ellison and Drew Fudenberg (1995) for a theoretical treatment of word-of-mouth communication. Matthew O. Jackson (2006) provides a recent survey of the literature on social networks.

² Indeed, recent theoretical work is beginning to explore the asset-pricing implications of information-sharing along social networks, while maintaining the assumption that such information-sharing is feasible (Paolo Colla and Antonio Mele 2005; Han N. Ozsoylev 2005).

³ To be more precise, such an honest exchange of information would seem particularly puzzling if it happens before either party has taken a position in the asset in question. Of course, if one manager already has a long position in a stock, he might wish to find a way to credibly communicate this to other managers, so that their buying pushes up the price of his position.

Yet if Silicon Valley firms do in fact compete intensely with one another, wouldn't any one executive be tempted either to withhold valuable information from his peers in other firms, or to actively mislead them, in the course of an informal conversation?

In this paper, I propose a theory of incentive-compatible information exchange among players who, as in the examples above, are in competition with one another. The theory has two key ingredients. First, information flows in both directions—from player A to player B and vice versa—during the course of a conversation, as players quite literally take turns bouncing ideas off of one another. This differs from the classic framework of Vincent P. Crawford and Joel Sobel (1982). Like I do, Crawford and Sobel pose the question of whether “cheap talk” can be credible in a situation where the two parties involved have partially conflicting interests. But in their model, one party is always the better-informed “sender,” and the other is always the less-well-informed “receiver,” so there is no scope for two-way communication.⁴

The second key ingredient is a strong form of complementarity in the information structure. In particular, player B can come up with the next idea in stage t of a conversation only if player A has disclosed his idea from stage $t-1$.⁵ This complementarity, combined with the two-way nature of the conversation, is what allows for incentive-compatible dialogue. If player A deviates from a truth-telling equilibrium at stage $t-1$ by lying to player B, he makes B less well-informed—which A prefers, since B is his competitor. At the same time, this also precludes B from coming up with a further idea at stage t which builds on A's previous insight, and which might have been passed back to A had the conversation continued in an honest fashion. This is a direct cost to player A. I show that if this latter effect is strong enough, it can be incentive-compatible for players to share their ideas truthfully with one another at all stages of a conversation.

The model can also be used to study the dynamics of information diffusion along social networks that encompass more than two players. Suppose that after A and B conclude their conversation, B can choose whether or not to tell a third player C what he has learned from the conversation, in the hopes that C may be able to further build on this idea. I show that B will do so only if his initial informational advantage over C is not too big—i.e., so long as B's prior conversation with A was not too productive.

This logic leads to a novel account of information diffusion in which relatively underdeveloped ideas can travel long distances and be shared by many agents, but in which more fully realized—and hence more valuable—ideas tend to remain localized among just a few agents. In other words, even if all agents in the economy are connected to the same network, incentive-compatibility considerations create an endogenous tendency for the most valuable ideas to stay confined to small groups.

The remainder of the paper is organized as follows. Section I models a single conversation among two players. Section II extends the basic framework to incorporate many players who are arrayed along a straight-line network, and who can have sequential conversations with their immediate neighbors. This extension allows one to ask how far a given set of ideas can travel along the network, and how the distance traveled relates to the quality of the ideas. Section III discusses the model's implications for Shiller's (2000) hypothesis, namely that word-of-mouth transmission of ideas can have significant consequences for asset prices. Section IV concludes.

⁴ The same can be said about a number of recent theoretical papers that study the process of verbal communication, e.g., Vijay Krishna and John Morgan (2004), and Mathias Dewatripont and Jean Tirole (2005). Although they substantially enrich the basic Crawford-Sobel (1982) model, they continue to maintain the sender-receiver dichotomy, so that information only ever flows in one direction.

⁵ Thomas F. Hellman and Enrico C. Perotti (2007) also emphasize the importance of complementarities in the production of ideas, though their focus is on a different set of issues. More generally, this “standing-on-the-shoulders” formulation is familiar in work on technological innovation (see, e.g., Philippe Aghion and Peter W. Howitt 1992).

I. The Basic Model: A Single Conversation

A. Information Structure

There are two players, A and B, engaged in a conversation. They take turns moving, so that A moves at times 1, 3, 5, etc., and B moves at times 2, 4, 6, etc. At time 1, player A has access to an initial idea \mathcal{X}_1 that he can choose to communicate to player B. This idea has two uses. First, anybody observing \mathcal{X}_1 can eventually “decode” it—i.e., can map it into a payoff-relevant signal $s_1(\mathcal{X}_1)$. This decoding is assumed to happen after the conversation has ended and A and B have gone their separate ways.

Second, if player B gains access to \mathcal{X}_1 , he has a probability p of coming up with another idea \mathcal{X}_2 at time 2 that builds on and refines \mathcal{X}_1 , and that can be mapped into a second payoff-relevant signal $s_2(\mathcal{X}_2)$. Crucially, if idea \mathcal{X}_1 is not truthfully revealed to player B, he can never come up with \mathcal{X}_2 . This assumption embeds a strong form of complementarity into the production function for ideas: a useful new idea can be produced only by an agent who has access to the prior idea.

The process can continue indefinitely, until one player fails to come up with an idea. In particular, if at any odd date t , player A has observed player B’s prior-period idea \mathcal{X}_{t-1} , he has a probability p of coming up with another idea \mathcal{X}_t , which he can then choose to communicate to B. Again, it is impossible for A to generate idea \mathcal{X}_t without having been exposed to a complete and truthful description of idea \mathcal{X}_{t-1} . And whoever observes \mathcal{X}_t can, after the conversation has ended, go off and generate the signal $s_t(\mathcal{X}_t)$.

Of course, a player in possession of an idea need not communicate it truthfully. In particular, if player A has come up with an idea \mathcal{X}_t at time t , he may either: not say anything to player B, effectively pretending that he was unable to come up with an idea; or pass along a bogus (i.e., informationally useless) version of the idea, \mathcal{X}_t^b . I assume that the former option is costless, but that the latter imposes on A an arbitrarily small positive cost of ε .⁶ Moreover, in the latter case, player B does not immediately recognize that he has been lied to—he takes the statement of \mathcal{X}_t^b at face value and attempts to come up with a further insight that builds on it. It is only after the conversation ends, when B attempts to decode the ideas that he has collected, that he learns that \mathcal{X}_t^b is useless: at that point, he is unable to turn \mathcal{X}_t^b into a payoff-relevant signal.

B. Payoffs

At the conclusion of the conversation, the players attempt to decode the ideas they have obtained. Let n_A represent the total number of signals that player A is able to produce (which equals the total number of nonbogus ideas that he has collected), and let n_B represent the total number of signals that player B is able to produce. To capture the notion that the players are competitors, A’s payoff should be an increasing function of n_A , and a decreasing function of n_B , and vice versa.

In order to generate payoff functions with transparent microfoundations, I consider a specific competitive setting with a simple market structure. As will become apparent, this setting connects more directly to the Silicon Valley production-technology example sketched in the introduction than to the fund-manager example.

Think of the signals generated by a conversation as elements of a recipe for building a product more efficiently. In particular, if player A has access to n_A signals, he can manufacture the product for a cost of $(1 - h(n_A))$, and symmetrically for player B. Here $h(n)$ is an increasing function

⁶ This is a technical assumption whose purpose will become clear shortly. One can motivate it by saying that it takes a little time and effort to come up with a good lie, or that lying is personally distasteful.

that captures the total cost savings associated with n signals. I will put more structure on the $h(n)$ function below, but for the moment I leave its general form unspecified, and just impose the conditions that $h(0) = 0$, and $h(\infty) < 1$.

I assume that player A and player B each face a unit mass of customers for the new product, and that all customers have a reservation value of one. Moreover, there is a fractional overlap of θ in A's and B's customer bases, with $0 < \theta < 1$. In other words, A has a monopoly on a fraction $(1 - \theta)$ of his customers, but must compete with B for the remaining fraction θ . I further assume that the products are otherwise undifferentiated, so that when competition between A and B does occur, it is à la Bertrand. In what follows, θ can be thought of as a proxy for the degree of competition between A and B.

Taken together, the assumptions imply that the payoffs to player A and player B are respectively given by:

$$(1) \quad U_A = (1 - \theta)h(n_A) + \theta \max\{0, h(n_A) - h(n_B)\},$$

$$(2) \quad U_B = (1 - \theta)h(n_B) + \theta \max\{0, h(n_B) - h(n_A)\}.$$

The first part of (1) corresponds to the fact that for a fraction $(1 - \theta)$ of his customers, A is a monopolist and charges the full reservation value of one; with a cost of $(1 - h(n_A))$, his profits per customer are thus $h(n_A)$. On the remaining fraction θ of his customers, where A and B overlap, Bertrand competition implies that A makes a profit only to the extent that his costs are strictly below those of B.

To give a concrete example, suppose that player A has been exposed to ideas \mathcal{X}_1 through \mathcal{X}_5 , while player B has been exposed only to ideas \mathcal{X}_1 through \mathcal{X}_4 ; this would happen if, after having observed \mathcal{X}_5 , A decided either to remain quiet, or to make a dishonest report of \mathcal{X}_5 to B. It follows that player A can generate the first five signals s_1 through s_5 , while player B can generate only the first four signals s_1 through s_4 . Thus we have $U_A = h(5) - \theta h(4)$, and $U_B = (1 - \theta)h(4)$.

C. Benchmark Case: No Complementarities in Idea Production

Before proceeding, it is useful to establish a benchmark result, namely that conversations can never arise absent complementarities in idea production of the sort described above. The following proposition summarizes the result.

PROPOSITION 1: *Suppose that a player's probability of having an idea \mathcal{X}_t at time t —and hence of being able to generate the signal $s_t(\mathcal{X}_t)$ —is fixed at p , and is independent of whether he has been given a truthful report of the prior idea \mathcal{X}_{t-1} . In this case, it can never be part of an equilibrium for players to truthfully exchange ideas with one another.*

PROOF:

Assume to the contrary that there does exist an equilibrium in which the players truthfully exchange ideas. Taking player B's honest reporting strategy as given, at any time t player A will wish to deviate from the proposed equilibrium, and mislead B with a dishonest report—i.e., A will want to report \mathcal{X}_t^b instead of \mathcal{X}_t . Such a report leaves B with one less signal than A. Moreover, since B does not know (until after the conversation has ended) that the report was dishonest, he cannot punish A for the deviation. Thus A does not experience any reduction in signals as a result of deviating; the only cost he bears is the ε cost of inventing the bogus report. With n_A unchanged, and n_B reduced to a value strictly lower than n_A (and with ε assumed to be arbitrarily small), equation (1) tells us that the deviation strictly increases player A's payoff U_A .

Intuitively, without complementarities in idea production, player A can only gain by deviating from the honest equilibrium and lying to player B about his ideas—such a lie leaves B less well informed, and taking B's truthful-reporting strategy as fixed, does not degrade the quality of the information that A expects to get back from B. The fact that it is impossible to sustain a cooperative outcome in this setting, in spite of an infinite horizon, distinguishes the no-complementarities version of the model from, e.g., an infinitely repeated prisoner's dilemma game. The key distinction is that in the prisoner's dilemma case, a deviation from the cooperative strategy is immediately observable and hence can be punished in a tit-for-tat fashion, while the same is not true here, since lies cannot be detected until after the conversation has concluded.

This logic also suggests why complementarities in idea production are necessary for an honest exchange of ideas. In a setting with complementarities, if A misleads B about his idea \mathcal{X}_t , A also stands to bear a cost. This is because without access to an honest report of \mathcal{X}_t , B can no longer come up with the next useful idea, which—assuming that B is playing an honest-reporting strategy—could otherwise have been bounced back to A.

D. Sustaining a Conversation: Necessary Conditions

I now return to the case where there are complementarities in idea production, and explore the conditions under which a cooperative conversational equilibrium can be sustained. To do so, I hypothesize that such an equilibrium exists, and then check that at each date t , neither player has an incentive to deviate from the proposed equilibrium.

In a fully cooperative equilibrium, each player truthfully reports any idea that he has to the other player. This continues until one player fails to come up with an idea, at which time he truthfully announces this as well, and the conversation concludes, with each player having access to the same total number of ideas.

Now suppose it is date t , and that player i has just come up with an idea \mathcal{X}_t . If he reports it honestly, and the game continues along the equilibrium path until somebody fails to come up with an idea, player i 's expected payoff, $EU_i(\text{continue}@t)$, is given by

$$(3) \quad EU_i(\text{continue}@t) = (1 - \theta)H(t),$$

where

$$(4) \quad H(t) = \sum_{i=0}^{\infty} p^i (1 - p) h(t + i).$$

In words, $H(t)$ is the expected value of the total cost reduction $h(\cdot)$ that both players will realize if player i shares \mathcal{X}_t , and they continue to play cooperatively from that point on. For example, with probability $(1 - p)$, there will be no further ideas after time t , so each player will wind up with t signals that have a total cost-cutting value of $h(t)$; with probability $p(1 - p)$, there will be exactly one further idea after t , so each player will wind up with $(t + 1)$ signals that have a total cost-cutting value of $h(t + 1)$; and so forth.

By contrast, suppose that player i considers deviating from the proposed cooperative equilibrium at time t . We can focus our attention on deviations that take the form of player i putting an end to the conversation by not reporting an idea \mathcal{X}_t that he is in possession of at time t .⁷ In this case, his payoff, $U_i(\text{stop}@t)$, is given by

⁷ In principle, a deviation could take the form of i either not reporting his idea to the other player, or inventing a bogus report. Unlike in the no-complementarities case of Proposition 1, however, now if i is going to deviate, it is a

$$(5) \quad U_i(\text{stop}@t) = h(t) - \theta h(t - 1).$$

This expression reflects the fact that if i deviates, he keeps idea \mathcal{X}_i to himself, and therefore winds up with one more signal than the other player (t signals versus $(t - 1)$). This allows him to earn not only a profit of $(1 - \theta)h(t)$ from the part of the market that he has all to himself, but also a profit of $\theta(h(t) - h(t - 1))$ on the part of the market that he has to compete for with the other player, since he now has a cost advantage.

In order for the cooperative equilibrium to hold together, we require that $EU_i(\text{continue}@t) \geq U_i(\text{stop}@t)$ for all values of t . Substituting in from equations (3) and (5), and doing some rearranging, this requirement can be expressed as

$$(6) \quad \frac{H(t) - h(t - 1)}{h(t) - h(t - 1)} \geq \frac{1}{(1 - \theta)}.$$

Thus we have:

PROPOSITION 2: *If condition (6) is satisfied for all values of t , it is possible to sustain a conversation in which each player truthfully reports every idea he has to the other player, until one of the two players fails to come up with a further idea.*

E. A Parametric Example: Geometric Decay in the Value of Successive Ideas

As stated, the necessary condition in (6) is not transparent, because it depends on the shape of the $h(n)$ function. To make things more intuitive, it is helpful to put some structure on this function. One particularly tractable choice is the following:

$$(7) \quad h(n) = (1 - \beta^n),$$

with $0 < \beta < 1$. This formulation corresponds to the assumptions that: as a conversation goes on, each new idea is of less value than the previous one; and that this decay is geometric in nature. In particular, the incremental cost savings associated with the n th signal is given by $\beta^{n-1}(1 - \beta)$. Thus the first signal lowers costs by $(1 - \beta)$, the second signal further lowers costs by $\beta(1 - \beta)$, and so forth.

With this functional form, condition (6) can be rewritten more specifically as

$$(8) \quad \frac{\left(\beta^{t-1} - \sum_{i=0}^{\infty} p^i (1 - p) \beta^{t+i} \right)}{(\beta^{t-1} - \beta^t)} \geq \frac{1}{(1 - \theta)}.$$

Condition (8) in turn can be greatly simplified to yield

dominant strategy for him to do so by not reporting anything. This is because in a world with complementarities, *any* deviation effectively ends the game from his perspective—once he stops feeding honest information to his counterpart, he can never get anything of value back. So there is no point in spending the ε cost of inventing a bogus report, when keeping quiet achieves the same outcome.

$$(9) \quad \frac{1}{(1 - p\beta)} \geq \frac{1}{(1 - \theta)},$$

or, alternatively,

$$(10) \quad p\beta \geq \theta.$$

The inequality in (10) captures in a simple way the three main factors that make it possible for conversations to be sustained. First, and most obviously, it helps if the competition parameter θ is small, so that one player suffers only a little bit when the other's information set is improved. Second, it also helps if p is large, which means that when one player truthfully reveals an idea to the other at time t , there is a relatively high likelihood that another idea will be bounced back to him at time $t + 1$; this makes continuing the conversation more attractive relative to deviating and cutting it short. And, finally, incentives for cooperation are also stronger when the decay parameter β is closer to one; this means that the marginal value of incremental ideas declines slowly, which also increases the appeal of continuing the conversation as opposed to cutting it off.

It is worth noting that the simple condition in (10) is independent of t . In other words, the incentive to continue a conversation does not depend on how far into it the players are. This is true even though the amount of information that is yet to be gained decreases with t . Intuitively, there are two offsetting effects. On the one hand, as time passes, the absolute appeal of continuation declines. On the other hand, so does the absolute appeal of deviating, and thereby staying one step ahead of one's competitor. Since all that matters is the ratio of these two quantities, they offset each other exactly, causing t to drop out of the comparison. This can be seen explicitly by looking at (8), where the numerator of the left-hand side (which reflects the appeal of continuation) and the denominator (which reflects the appeal of deviating) are both proportional to β^{t-1} , causing β^{t-1} to drop out when the expression is simplified.

II. Sequential Conversations: Information Diffusion on a Social Network

A. Extending the Basic Model

I now augment the model so that there are multiple players arrayed along a simple social network, namely a straight line. I assume that each player interacts only with his immediate neighbors, and that this interaction happens sequentially, starting with players A and B. Thus, after A and B conclude their conversation, B must choose whether to initiate a conversation with player C. If this happens, when B and C conclude their conversation, C can choose whether to start one with player D, and so on.

In order to generate the players' payoff functions, I consider the natural extension of the market structure from the two-player case. Specifically, every player in the interior of the network (i.e., every player but A) now has a customer base that overlaps partially with his neighbors on either side. For example, of the unit mass of customers facing player B, a fraction θ overlaps with player A, another fraction θ overlaps with player C, and a fraction $(1 - 2\theta)$ are customers who are exclusive to B—i.e., customers over whom B has monopoly power. (Note that I must now assume that $\theta < 1/2$, which is innocuous.)

It follows that player A's payoff function remains as in equation (1), while the payoff function for player B is modified to

$$(11) \quad U_B = (1 - 2\theta)h(n_B) + \theta \max\{0, h(n_B) - h(n_A)\} + \theta \max\{0, h(n_B) - h(n_C)\}.$$

An analogous expression holds for the payoffs of each of the players who come after B.

B. Necessary Conditions for a Second Conversation to Get Started

Assume for the moment that the first conversation—that between A and B—is itself incentive-compatible. As we have seen, this conversation will go on until one of the two players is unable to come up with a further refinement. At this point, A and B part ways. Let k denote the random number of signals that A and B are able to generate after their conversation concludes. That is, both A and B have access to idea \mathcal{X}_k at the end of their conversation, and can therefore both manufacture signals s_1 through s_k .

Next, B has to decide whether to turn around and share idea \mathcal{X}_k with player C. As before, the potential advantage of doing so is that player C will, with probability p , be able to come up with a further refinement \mathcal{X}_{k+1} , and that the conversation between B and C will continue on for several more stages from there. The disadvantage to B of starting a conversation with C is that by giving C access to idea \mathcal{X}_k , B effectively gives C *all* of the first k signals s_1 through s_k —i.e., B repeats everything of value that was learned during the A-B conversation. The implicit assumption here is that the refined idea \mathcal{X}_k also embodies all of the cumulative knowledge in the previous-stage ideas \mathcal{X}_1 through \mathcal{X}_{k-1} .

Player B's incentive to initiate a conversation with C is therefore lower than A's original incentive to initiate a conversation with B. This is because when A first speaks to B, he gives away a brand-new idea that is equivalent to just one signal, in the hopes of getting one more signal in return. By contrast, when B first speaks to C, he gives away a more fully developed idea that is equivalent to k signals, again in the hopes of getting one more signal in return.

To make this precise, observe that at the conclusion of the A-B conversation, with both A and B in possession of k signals, we can rewrite B's payoff as

$$(12) \quad U_B(k) = (1 - 2\theta)h(n_B) + \theta \max\{0, h(n_B) - h(k)\} + \theta \max\{0, h(n_B) - h(n_C)\} \\ = (1 - \theta)h(n_B) - \theta h(k) + \theta \max\{0, h(n_B) - h(n_C)\}.$$

Based on (12), player B reasons as follows. If he reveals his information to player C, and the conversation continues from there, player B's expected payoff from initiating the B-C conversation is given by⁸

$$(13) \quad EU_B(\text{initiate@BC}) = (1 - \theta)H(k) - \theta h(k).$$

Note that conditional on completing a conversation with player C, player B does not care if C then goes ahead and improves his information set via a subsequent conversation with player D. This is because once B and C have access to the same number of signals, B can never earn a profit from those customers they have in common; this does not change if C eventually becomes even better informed than B.

If, instead, player B decides not to speak to player C, thereby eliminating the potential for any further conversations, he walks away with a payoff of

⁸ One still has to check the incentive-compatibility conditions that ensure that once B and C have started a conversation, they will both want to continue it until they run out of ideas. However, it is easy to show that these conditions are implied by (6). Also, note that once a conversation is started, C will always have a stronger incentive to continue it than B. This is because C has the option value associated with the possibility of moving on to a further conversation with D.

$$(14) \quad U_B(\text{withdraw@BC}) = (1 - \theta)h(k).$$

Therefore, a necessary condition for player B to initiate a conversation with C is

$$(15) \quad \frac{H(k)}{h(k)} \geq \frac{1}{1 - \theta}.$$

Comparing (15) to (6), it is apparent that, for $k = t$, (15) implies (6), so that it is harder to satisfy (15), all else equal. This formalizes the intuition stated above: at any point in time t , it is easier to get a player to continue an existing conversation than it is to get a player with an equally well-developed idea to initiate a new conversation: in the former case, the player is being asked to give away only one incremental signal, while in the latter, he is being asked to start things off by giving away k signals.

C. When Does Information Diffusion Come to a Stop?

For concreteness, the remainder of the analysis in the paper focuses on the parametric specification described above, where the value of new signals decays geometrically. With this specification, condition (15) boils down to

$$(16) \quad \frac{1 - (\beta^k(1 - p)/(1 - p\beta))}{(1 - \beta^k)} \geq \frac{1}{(1 - \theta)}.$$

Observe that the previous necessary condition for continuing a conversation in the two-player model, given by (10), is just a special case of (16) with $k = 1$. In other words, if player B winds up his conversation with A with just one signal, his decision of whether to initiate a new conversation with C is identical to his decision of whether to continue an existing conversation in the two-player version of the model, since either choice amounts to giving away one signal in exchange for the same expected informational return.

If (16) is in fact satisfied for a given realization of k , then the conversation between B and C gets off the ground. This conversation then continues stochastically until one of the two is unable to generate a further refinement. Suppose that this happens at the point where B and C can each generate k' signals, with $k' \geq k$. Now C faces the exact same problem in deciding whether to initiate a conversation with D, so he applies the same necessary condition as B did previously, except that k is replaced in (16) with k' .

Using this sequential logic, we can describe the dynamics of information diffusion in the multiplayer version of the model. Note that the left-hand side of (16) is strictly decreasing in k . Now define k^* as that value of k for which (16) holds with equality:

$$(17) \quad k^* = \frac{\log(1 - p\beta) + \log(\theta) - \log(\theta(1 - p) + p(1 - \beta))}{\log(\beta)}.$$

The following proposition describes the outcome for the geometric-decay case:⁹

⁹ It is possible to show that an analog to Proposition 3 holds for a broader class of payoff functions, $h(n)$. Specifically, suppose that $h(n)$ is a strictly increasing and strictly concave function such that $h(0) = 0$, and $\lim_{n \rightarrow \infty} h(n) < 1$, and that the general condition in (6) for sustaining a conversation is satisfied for all t . One can then show that there exists a k^* with $1 \leq k^* < \infty$ such that the proposition continues to apply.

PROPOSITION 3: *Consider two adjacent players J and K in the network, and assume that (10) is satisfied, so that once initiated, conversations can always be sustained. Assume further that J and K conclude a conversation with each having the ability to produce k_{JK} signals. If $k_{JK} \leq k^*$, player K then initiates a new conversation with the next player L in the network, and the process of information diffusion continues forward. If $k_{JK} > k^*$, no further conversations are initiated.*

Example 1: Suppose that $\beta = 0.95$, $\theta = 0.1$, and $p = 0.5$. Then, according to (17), $k^* = 6.95$. Thus, if a conversation between J and K ends with each able to produce six or fewer signals, K initiates a new conversation with L. If a conversation between J and K ends with each able to produce seven or more signals, no further conversations are initiated.

D. Implications for Number of Participants, Quality of Ideas, and Conformity

With Proposition 3 in hand, it is possible to establish a number of properties of the sequential model. The statements that follow center on two items of interest. First, the random variable c is defined as the number of conversations that occur in a given play of the game *subsequent* to the initial A-B conversation (which, under the assumption that (10) is satisfied, always takes place). Second, the random variable k , which we have already encountered, is the duration of the initial A-B conversation. One way to think about k is that it measures the quality of the idea that emerges from this first conversation. Alternatively, k can be interpreted as a measure of the degree of pairwise conformity among participants in the conversational chain. Specifically, for any given play of the game, k is equal to the number of signals that the initial player A holds in common with any other player J who has participated in the chain. Thus k is a summary statistic for the extent to which A's and J's production technologies (or forecasts, depending on the setting) resemble one another.

As described in the introduction, one of the main insights of this paper is that when ideas travel farther, they tend to be of lower quality, and hence to induce a weaker degree of conformity among those who have been exposed to them. There are two ways that this insight can be expressed in the language of the model. The first, and perhaps most natural, way to do so is to fix the parameters p , β , and θ , and to make statements about the correlation between the ex post realizations of c and k . The Appendix (available at <http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.5.2150>) establishes the following:

PROPOSITION 4: *Let ρ_{ck} denote the correlation between the random variables c and k . For any values of the parameters p , β , and θ such that (10) holds, we have that $\rho_{ck} < 0$.*

The logic behind the proposition is straightforward. Suppose that the initial conversation between A and B is blessed by good luck, and continues for many rounds, so that the number of signals k is relatively high. In this case, k is closer to k^* , which represents the cutoff point beyond which no new conversations are started; it is therefore more likely that the whole process will soon come to a stop. Stated a bit differently, the expected number of subsequent conversations conditional on k , $E(c|k)$, is declining in k .

An alternative way to think about the relationships between the number of conversations, idea quality, and conformity is to compute the ex ante expected values of c and k , and then to see how these expected values vary as we change one of the exogenous parameters, say p . The Appendix shows that there are simple closed-form expressions for the expected values of c and k :

PROPOSITION 5: *Assume that (10) holds, and define $I(k^*)$ as the greatest integer that is less than or equal to k^* . Then we have: $E(c) = I(k^*)(1 - p)/p$, and $E(k) = 1/(1 - p)$.*

The comparative statics that follow from Proposition 5 are for the most part intuitively obvious. First, it is immediate that $dE(k)/dp > 0$. Since one can interpret p as a proxy for the talent of the players in a given network—more talented players are less likely to draw a blank at any given point in a conversation—this just says that both the expected quality of the idea emerging from the initial A-B conversation, as well as conformity, are increasing in talent. Second, it can be shown that $E(c)$ is weakly decreasing in θ , which means that the expected number of participants in a conversational chain increases when players are less competitive with one another.

The one comparative static that is a little tricky is $dE(c)/dp$. As it turns out, $E(c)$ is not monotonic in p , because of the discreteness associated with the integer-valued function $I(k^*)$. However, if we get rid of this discreteness issue, it is possible to make a clearer statement. Specifically, let $\Phi\{E(c)\}$ denote the upper bound on $E(c)$ defined by

$$(18) \quad \Phi\{E(c)\} = k^*(1 - p)/p.$$


It can then be shown that $d\Phi\{E(c)\}/dp < 0$, as illustrated in Figure 1, which plots both $E(c)$ and $\Phi\{E(c)\}$ as functions of p for $\beta = 0.95$ and $\theta = 0.1$. In other words, as p increases, the upper bound on the expected number of conversations falls. More precisely, an increase in p has two competing effects on $\Phi\{E(c)\}$. First, the expected length of any conversation goes up, and so it becomes more likely that any given cutoff point k^* is reached after a smaller number of conversations. Second, however, the cutoff point k^* is itself an increasing function of p , since the desirability of initiating new conversations rises with talent. As it turns out, the former effect dominates the latter, so that $d\Phi\{E(c)\}/dp < 0$. This is established formally in the Appendix.

Thus there is another sense in which we can think of the number of conversations as being negatively related to both idea quality and conformity. As we increase the talent p of the players, the expected quality of the idea generated in the initial A-B conversation goes up, as does conformity; this is just the statement that $dE(k)/dp > 0$. And the same increase in talent p also lowers the upper bound on the expected number of conversations, as we have just seen. So if there is variation in p across social networks, it will trace out the aforementioned negative relationships. For example, if one social network is composed of highly talented managers, and another network is composed of lower-ability managers, the model predicts that conversational chains would be on average shorter in the former case, but that they would tend to transmit higher-quality ideas, and to induce more conformity among those managers that end up participating in the conversations.

Example 2: First, suppose that, as in the previous example, $\beta = 0.95$, $\theta = 0.1$, and $p = 0.5$. Straightforward calculation based on Proposition 5 yields: $E(c) = 6$, $\Phi\{E(c)\} = 6.95$, and $E(k) = 2$. If we keep everything else the same but set $p = 0.8$, we obtain instead: $E(c) = 4.25$, $\Phi\{E(c)\} = 4.47$, and $E(k) = 5$.

III. Conversations and Asset Prices

I have framed the discussion of the model in terms of a specific example, one in which “ideas” represent cost-saving innovations, and in which the “competitors” are rival producers. Of course, this is not the only potentially relevant setting. As noted in the introduction, Shiller (2000) hypothesizes that word-of-mouth transmission of ideas can play an important role in generating asset-price volatility. In this context, “ideas” presumably represent signals about the expected returns to particular investment strategies, and the “competitors” are professional investors.

FIGURE 1. $E(c)$ AND $\Phi\{E(c)\}$ AS FUNCTIONS OF p

Note: The figure plots both $E(c) = I(k^*)(1 - p)/p$ and its upper bound, $\Phi\{E(c)\} = k^*(1 - p)/p$, as functions of p for $\beta = 0.95$ and $\theta = 0.1$.

Applied this way, the model can help to rationalize the empirical findings of Hong, Kubik, and Stein (2005) and Cohen, Frazzini, and Malloy (2007). These papers show that there are clusters of correlation in the trades of mutual-fund managers who are either located in the same city or who went to college together—patterns that strongly suggest word-of-mouth diffusion of investment ideas. And, notably, these effects show up among agents who have clear incentives to care about relative performance.

Beyond such patterns in trading behavior, can the mechanism in the model also generate interesting price effects of the sort envisioned by Shiller (2000)? Here the answer is less clear-cut. Note that a given idea will have more of an impact on stock prices to the extent that two conditions are satisfied: the idea reaches a large number of investors; and its information content is substantial, so that each investor who comes into contact with the idea revises his expectations by a meaningful amount. However, a central message from the model is that there can be a fundamental tension between these two conditions: those ideas that reach the largest numbers of investors are precisely those that tend to have relatively little real information content.¹⁰ Indeed, in the limit, such far-reaching ideas may not be much more substantial than gossip. This in turn would seem to imply an endogenous upper bound on the magnitude of any word-of-mouth effect.

One caveat is that the upper bound argument relies on the premise that investors are rational, so that the stock-price impact associated with an idea is proportional to its true informational content. In reality, however, it may be that information obtained via face-to-face interaction is excessively salient and compelling to investors, as compared to, say, relatively dry earnings

¹⁰ To map the model more literally into a financial-markets setting, suppose that each player wants to forecast the payoff V on a security, which can be written as an infinite sum, $V = s_1 + \dots + s_\infty$, where s_i is a normal random variable with a variance of β^i . This formulation mirrors the setup in the text: the marginal value of each successive signal—here in terms of the fraction of the variance of V that it explains—declines geometrically. And a low-information-content, far-traveling idea is one that is comprised of just a small number of signals, and hence explains only a small fraction of the overall variance of V .

releases. If so, even low-quality rumors might have significant consequences for prices, which clearly helps Shiller's story.

IV. Conclusions

The transmission of information by word-of-mouth occurs in a wide range of settings, including some in which agents have competing interests. In the face of such competition, an information structure with complementarities can be the glue that holds conversations together: it gives each participant in a conversation an incentive to honestly disclose a given idea, in the hopes that his counterpart will be able to take it one step further, and will then bounce the more fully developed idea right back to him.

While an information structure with complementarities makes conversations among competitors feasible, it does not remove all barriers to communication. Thus, while ideas can travel across networks of competitors who are connected to one another, they do not necessarily travel very far. Indeed, the more successfully an idea develops in early-stage conversations, the more likely it is to remain localized among the handful of players who were its originators. In contrast, low-quality ideas—those not much more informative than gossip—may ultimately be very widely diffused.

REFERENCES

- Aghion, Philippe, and Peter Howitt.** 1992. "A Model of Growth through Creative Destruction." *Econometrica*, 60(2): 323–51.
- Cohen, Lauren H., Andrea Frazzini, and Christopher J. Malloy.** 2007. "The Small World of Investing: Board Connections and Mutual Fund Returns." Unpublished.
- Colla, Paolo, and Antonio Mele.** 2005. "Information Linkages and Correlated Trading." Unpublished.
- Crawford, Vincent P., and Joel Sobel.** 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431–51.
- Dewatripont, Mathias, and Jean Tirole.** 2005. "Modes of Communication." *Journal of Political Economy*, 113(6): 1217–38.
- Ellison, Glenn, and Drew Fudenberg.** 1995. "Word-of-Mouth Communication and Social Learning." *Quarterly Journal of Economics*, 110(1): 93–125.
- Hellmann, Thomas F., and Enrico C. Perotti.** 2007. "The Circulation of Ideas in Firms and Markets." Unpublished.
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein.** 2005. "Thy Neighbor's Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers." *Journal of Finance*, 60(6): 2801–24.
- Jackson, Matthew O.** 2006. "The Economics of Social Networks." In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Volume 1*, eds. Richard Blundell, Whitney Newey, and Torsten Persson, 1–56. New York: Cambridge University Press.
- Krishna, Vijay, and John Morgan.** 2004. "The Art of Conversation: Eliciting Information from Experts through Multi-Stage Communication." *Journal of Economic Theory*, 117(2): 147–79.
- Ozsoylev, Han N.** 2005. "Asset Pricing Implications of Social Networks." Unpublished.
- Saxenian, Anna Lee.** 1994. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.
- Shiller, Robert J.** 2000. *Irrational Exuberance*. Princeton, NJ: Princeton University Press.